



HHS Public Access

Author manuscript

Med Image Anal. Author manuscript; available in PMC 2022 February 01.

Published in final edited form as:

Med Image Anal. 2021 February ; 68: 101889. doi:10.1016/j.media.2020.101889.

Capsules for Biomedical Image Segmentation

Rodney LaLonde^a, Ziyue Xu^b, Ismail Irmakci^c, Sanjay Jain^d, Ulas Bagci^{a,*}

^aCenter for Research in Computer Vision (CRCV), University of Central Florida (UCF), Orlando, FL

^bNvidia, Bethesda, MD

^cEge University, Izmir, Turkey

^dJohns Hopkins University, Baltimore, MD

Abstract

Our work expands the use of capsule networks to the task of object segmentation for the first time in the literature. This is made possible via the introduction of *locally-constrained routing* and *transformation matrix sharing*, which reduces the parameter/memory burden and allows for the segmentation of objects at large resolutions. To compensate for the loss of global information in constraining the routing, we propose the concept of “*deconvolutional*” capsules to create a deep encoder-decoder style network, called *SegCaps*. We extend the masked reconstruction regularization to the task of segmentation and perform thorough ablation experiments on each component of our method. The proposed convolutional-deconvolutional capsule network, *SegCaps*, shows state-of-the-art results while using a fraction of the parameters of popular segmentation networks. To validate our proposed method, we perform experiments segmenting pathological lungs from clinical and pre-clinical thoracic computed tomography (CT) scans and segmenting muscle and adipose (fat) tissue from magnetic resonance imaging (MRI) scans of human subjects’ thighs. Notably, our experiments in lung segmentation represent the largest-scale study in pathological lung segmentation in the literature, where we conduct experiments across five extremely challenging datasets, containing both clinical and pre-clinical subjects, and nearly 2000 computed-tomography scans. Our newly developed segmentation platform outperforms other methods across all datasets while utilizing less than 5% of the parameters in the popular *U-Net* for biomedical image segmentation. Further, we demonstrate capsules’ ability to generalize to unseen handling of rotations/reflections on natural images.

Graphical Abstract

*Corresponding author: Tel.: +1-407-823-1047; fax: +1-407-823-0594; bagci@crcv.ucf.edu (Ulas Bagci).

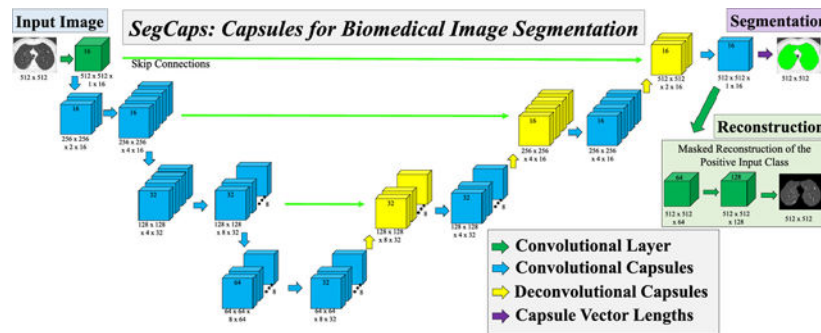
Authors’ Contributions

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



Keywords

Capsule Network Lung Segmentation Pre-Clinical Imaging

1. Introduction

The task of segmenting objects from images can be formulated as a joint object recognition and delineation problem. The goal in *recognition* is to locate an object's presence in an image, whereas delineation attempts to draw the object's spatial extent and composition (Bagci et al., 2012). Solving these tasks jointly (or sequentially) results in partitions of non-overlapping, connected regions, homogeneous with respect to some signal characteristics. Object segmentation is an inherently difficult task; apart from recognizing the object, we also have to label that object at the pixel level, which is an ill-posed problem.

Segmentation is of significant importance in biomedical image analysis, aiding systems focused on localizing pathologies (A. El-Baz et al., 2006), tracking disease progression (Xu et al., 2019), characterizing anatomical structure and defects (Frag et al., 2005), and many more (Elnakib et al., 2011). Due to its significance in many applications, segmentation is an essential part of most computer-aided diagnosis (CAD) systems, where the functionality of such systems can be heavily dependent on the accuracy of the segmentation module. Medical image segmentation brings its own set of unique challenges. Many anatomical structures vary significantly across individuals, with the presence of pathologies adding an additional layer of variation and complexity. Further, scanner artifacts and other noise can make the segmentation suboptimal. Recently convolutional neural network (CNN) methodologies have dominated the segmentation field, both in computer vision and medical image segmentation, most notably *U-Net for biomedical image segmentation* (Ronneberger et al., 2015), due to their remarkable predictive performance.

1.1. Drawbacks of CNNs and How Capsules Solve Them

The CNNs, despite showing remarkable flexibility and performance in a wide range of computer vision tasks, do come with their own set of flaws. Due to the scalar and additive nature of neurons in CNNs, neurons at any given layer of a network are ambivalent to the spatial relationships of neurons within their kernel of the previous layer, and thus within their effective receptive field of the given input. Feature maps in CNNs only contain scalar values, whether or not a given feature is present at each scalar location. These maps are

created from one layer to the next by multiplying each previous layer feature map by a set of kernel, then summing their activations to designate the presence/absence of the next higher-level feature. Since CNNs only have the ability to add presence activations within local kernels, higher-level neurons can only identify features within their effective receptive fields, but they cannot describe those feature in any way (*e.g.* precise location information, pose, deformation, etc.) To address this significant shortcoming, Sabour et al. (2017) introduced the idea of *capsule networks*, where information at the neuron level is stored as vectors, rather than scalars. These vectors contain information about:

1. spatial orientation and location information,
2. magnitude/prevalence, and
3. other attributes of the extracted feature

represented by each capsule type of that layer. This solves the previous issue of precise spatial localization in CNNs because capsule vectors can now additionally rate-code the exact position within the effective receptive field of the already place-coded vectors within the dimensions of those vectors. These sets of neurons, henceforth referred to as capsule types, are then “routed” to capsules in the next layer via a *dynamic routing algorithm* which takes into account the agreement between these capsule vectors, thus forming meaningful part-to-whole relationships not found in standard CNNs.

The overall goal of this study is to extend capsule networks and the dynamic routing algorithm to accomplish the task of object segmentation for the first time in the literature. We hypothesize that capsules can be used effectively for object segmentation with high accuracy and heightened efficiency compared to the state-of-the-art segmentation methods. To show the efficacy of the capsules for object segmentation, we choose a challenging application of pathological lung segmentation from computed tomography (CT) scans, where we have analyzed the largest-scale study of data obtained from both clinical and pre-clinical subjects, comprising nearly 2000 CT scans across five datasets and muscle and adipose (fat) tissue segmentation from magnetic resonance imaging (MRI) scans of three different contrasts obtained from a cohort of 50 patients (150 scans). We chose pathological lung segmentation for its obvious life-saving potential and unique challenges such as high intra-class variation, noise, artifacts and abnormalities, and other reasons discussed in Section 2. The additional experiments on muscle and adipose (fat) tissue segmentation compliment these first experiments both in the modality of the imaging technology used (MRI vs. CT) and anatomical structure. To further demonstrate the general applicability of our methods, we also provide proof-of-concept results for rotations/reflections on standard computer vision images showing, the ability of a capsule-based segmentation network to generalize to unseen poses of objects, a strong motivation for choosing capsule networks over CNNs in segmentation applications.

1.2. Building Blocks of Capsules for Segmentation

Performing object segmentation with a capsule-based network is extremely difficult due to the added computational cost of storing and routing vector representations, rather than scalars. The original capsule network architecture and dynamic routing algorithm is

extremely computationally expensive, both in terms of memory and run-time. Additional intermediate representations are needed to store the output of “child” capsules in a given layer while the dynamic routing algorithm determines the coefficients by which these children are routed to the “parent” capsules in the next layer. This dynamic routing takes place between every parent and every possible child. One can think of the additional memory space required as a multiplicative increase of the batch size at a given layer by the number of capsule types at that layer. The number of parameters required quickly swells beyond control as well, even for trivially small inputs such as MNIST and CIFAR10. For example, given a set of 32 capsule types with 6×6 , 8D-capsules per type being routed to 10×1 , 16D-capsules (as is the case in CapsNet), the number of parameters for this layer alone is $10 \times (6 \times 6 \times 32) \times 16 \times 8 = 1,474,560$ parameters. This one layer contains, coincidentally, roughly the same number of parameters as our entire proposed deep convolutional-deconvolutional capsule network with locally-constrained dynamic routing which itself operates on up to 512×512 pixel inputs. To scale the original CapsNet up to 512×512 , without these novelties, would require $512 \times 512 \times 32 \times 8 \times 512 \times 512 \times 10 \times 16 = 2814749767106560$ parameters and hence over 10 million GB of memory to store the parameters alone.

We solve this memory burden and parameter explosion by extending the idea of convolutional capsules (primary capsules in Sabour et al. (2017) are technically convolutional capsules without any routing) and rewriting the dynamic routing algorithm in two key ways. First, children are only routed to parents within a defined spatially-local kernel. Second, transformation matrices are shared for each member of the grid within a capsule type but are not shared across capsule types. To compensate for the loss of global connectivity with the locally-constrained routing, we extend capsule networks by proposing “deconvolutional” capsules which operates using transposed convolutions, routed by the proposed locally-constrained routing. These innovations allow us to still learn a diverse set of different capsule types while dramatically reducing the number of parameters in the network, addressing the memory burden. Also, with the proposed deep convolutional-deconvolutional architecture, we retain near-global contextual information and produce state-of-the-art results for our given applications. Our proposed *SegCaps* architecture is illustrated in Figure 2. As a comparative baseline, we also implement a simple three-layer capsule structure, more closely following that of the original capsule implementation, shown in Figure 1.

1.3. Summary of Our Contributions

The novelty of this paper can be summarized as follows:

1. Our proposed *SegCaps* is the first use of a capsule network architecture for object segmentation in literature.
2. We propose two technical modifications to the original dynamic routing algorithm where (i) children are only routed to parents within a defined spatially-local window and (ii) transformation matrices are shared for each member of the grid within a capsule type. These modifications, combined with convolutional capsules, allow us to operate on large images sizes (up to 512×512 pixels) for

the first time in literature, where previous capsule architectures typically do not exceed inputs of 32×32 pixels in size.

3. We introduce the concept of "deconvolutional" capsules and create a novel deep convolutional-deconvolutional capsule architecture, far deeper than the original three-layer capsule network, implement a three-layer convolutional capsule network baseline using our locally-constrained routing to provide a comparison with our *SegCaps* architecture, and extend the masked reconstruction of the target class as a method for regularization to the problem of segmentation as described in Section 3.
4. We validate the efficacy of *SegCaps* on the largest-scale study for pathological lung segmentation in the literature, comprising five datasets from both clinical and pre-clinical subjects with nearly 2000 total CT scans, and 150 MRI scans at three different contrasts for thigh muscle and adipose (fat) tissue segmentation. For lung segmentation, our proposed method produces improved results in terms of dice coefficient and Hausdorff distance (HD), when compared with state-of-the-art methods *U-Net* (Ronneberger et al., 2015), *Tiramisu* (Jégou et al., 2017), and *P-HNN* (Harrison et al., 2017), while dramatically reducing the number of parameters needed to achieve this performance. The proposed *SegCaps* architecture contains 4.6% of the parameters of *U-Net*, 9.5% of *P-HNN*, and 14.9% of *Tiramisu*. Thorough ablation studies are also performed to analyze the contribution and effect of several experimental settings in our proposed model. In particular, there is no other study to conduct fully-automated deep learning based pre-clinical image segmentation due to the extreme levels of variation in both anatomy and pathology present in animal subjects, the large number of high-resolution slices per scan with typically high levels of noise and scanner artifacts, as well as the sheer difficulty in even establishing ground-truth labels compared to human-subject scans.

1.4. Preliminary Non-Archival SegCaps Study

In a preliminary non-archival study, we introduced the first-ever capsule-based segmentation network in the literature, which we named *SegCaps* (LaLonde and Bagci, 2018). This initial study demonstrated the ability of SegCaps to perform on-par with state-of-the-art CNNs on the task of pathological lung segmentation from a dataset of CT scans. However, the study was extremely limited in a number of regards: 1) The ground-truth for the dataset was provided by an automated algorithm, raising concerns of similar biases/errors in the proposed algorithm existing within the ground-truth; 2) Results were only presented in terms of Dice scores which do not accurately capture the quality of the segmentation boundaries for the large lung fields; 3) The amount of pathology present in LUNA challenge dataset was fairly limited, targeted at lung nodule detection, rather than segmentation. In this journal extension of our non-archived work, we significantly extend the applications of SegCaps to better demonstrate its generalization abilities and versatility, performing segmentation in both CT and MRI on a wide range of anatomical structures. We obtained radiologist provided annotations for five large-scale pathological lung datasets from both clinical and preclinical subjects covering a wide range of pathologies, as well as dataset of three different

MRI contrasts for segmenting muscle and adipose (fat) tissue. We also perform a set of experiments to show a capsule-based segmentation network can better handle changes to viewpoint than a CNN-based approach.

2. Background and Related Works

Object segmentation in the medical imaging and computer vision communities has remained an interesting and challenging problem over the past several decades. Early attempts in automated object segmentation were analogous to the if-then-else expert systems of that period, where the compound and sequential application of low-level pixel processing and mathematical models were used to build-up complex rule-based systems of analysis (Horowitz and Pavlidis, 1974; Rosenfeld and Kak, 1982). In computer vision fields, superpixels and various sets of feature extractors such as scale-invariant feature transform (SIFT) (Lowe, 1999) or histogram of oriented gradients (HOG) (Dalal and Triggs, 2005) were used to construct these spaces. Specifically in medical imaging, methods such as level sets (Vese and Chan, 2002), fuzzy connectedness (Udupa and Samarasekera, 1996), graph-based (Felzenszwalb and Huttenlocher, 2004), random walk (Grady, 2006), and atlas-based algorithms (Pham et al., 2000) have been utilized in different application settings. Over time, the community came to favor supervised machine learning techniques, where algorithms were developed using training data to teach systems the optimal decision boundaries in a constructed high-dimensional feature space.

In the last few years, deep learning methods, in particular convolutional neural networks (CNNs), have become the state-of-the-art for various image analysis tasks (Ren et al., 2015; He et al., 2016, 2017; Huang et al., 2017; Hu et al., 2018). Specifically related to the object segmentation problem, *U-Net* (Ronneberger et al., 2015), Fully Convolutional Networks (FCN) (Long et al., 2015), and other encoder-decoder style CNNs have become the desired models for various medical image segmentation tasks. Most recent attempts in the computer vision and medical imaging literature utilize the extension of these methods to address the segmentation problem (Zhao et al., 2017; Chen et al., 2018b; Yang et al., 2018).

2.1. CNN-Based Segmentation

The object segmentation literature is vast, both before and in the deep learning era. Herein, we only summarize the most popular deep learning-based segmentation algorithms. Based on FCN (Long et al., 2015) for semantic segmentation, *U-Net* (Ronneberger et al., 2015) introduced an alternative CNN-based pixel label prediction algorithm which forms the backbone of many deep learning-based segmentation methods in medical imaging today. Following this, many subsequent works follow this encoder-decoder structure, experimenting with dense connections, skip connections, residual blocks, and other types of architectural additions to improve segmentation accuracy for particular imaging applications. For instance, a recent example by Jégou et al. (2017) combines a *U-Net*-like structure with the very successful DenseNet (Huang et al., 2017) architecture, creating a densely connected *U-Net* structure, called *Tiramisu*. Other successful frameworks for segmentation and their specific innovations are the following.

SegNet (Badrinarayanan et al., 2017) attempts to improve the upsampling process by performing “unpooling”, capturing the pooling indices from the max pooling layers in the encoder to more accurately place features in the decoder feature maps. Although the encoder-decoder structure is specifically designed to capture global context information, several methods attempt to further improve this global context in different ways. RefineNet (Lin et al., 2017) fuses features from multiple resolutions through adding residual connections and chained residual pooling to create a large cascaded encoder-decoder structure. PSPNet (Zhao et al., 2017) introduces a pyramid pooling module by pooling at different kernel sizes and concatenating back to the features maps. Large Kernel Matters (Peng et al., 2017) uses large $1 \times 15 + 15 \times 1$ and $15 \times 1 + 1 \times 15$ global convolution networks. ClusterNet (LaLonde et al., 2018) combines two fully-convolutional networks, one to capture global and one for local information, to segment specifically a large number of densely packed tiny objects, normally lost in networks with pooling. DeepLab (Chen et al., 2018a) utilizes an atrous spatial pyramid pooling (ASPP) unit to better capture image context from multiple scales. The latest version of DeepLab (v3+) (Chen et al., 2018b) follows a very similar structure to *U-Net* with the addition of an ASPP for image context and depthwise separable convolutions for efficiency.

2.2. Segmentation in Biomedical Imaging

As mentioned in Section 1, segmentation is of critical importance as a first stage in many biomedical imaging applications. Though well motivated, performing segmentation within biomedical imaging introduces a number of unique challenges, including handling many different imaging modalities anatomical structures, and potential deformities/abnormalities caused by a wide range of reasons. Further, imaging data across different applications can be 2D, 3D, and even 4D, requiring unique considerations. Multi-view networks, such as (Mortazi et al., 2017a), remain a popular approach to handling imaging data with more than two dimensions. 3D networks such as 3D U-Net (Çiçek et al., 2016) and V-Net (Milletari et al., 2016) have also gained recent popularity based off the highly successful U-Net (Ronneberger et al., 2015). Nonetheless, due to a combination of limited GPU memory and the desire to exploit existing pretrained models, majority of the literature uses 2D network and analyzes 3D data in a slice-wise manner. In this study, we focus on two of the most commonly investigated imaging modalities, namely CT and MRI, and detail the specific related works to those applications in the following paragraphs.

2.3. Pathological Lung Segmentation from CT

Anatomy and pathology segmentation have been central to the most medical imaging applications. Despite its importance, accurate segmentation of pathological lungs from CT scans remains extremely challenging due to a wide spectrum of lung abnormalities such as consolidations, ground glass opacities, fibrosis, honeycombing, tree-in-buds, and nodules. Specifically developed for pathological lung segmentation, Mansoor et al. (2014) created a two-stage approach based on fuzzy connectedness and texture features, incorporating anatomical information by segmenting the rib-cage. Most recently, Harrison et al. (2017) developed *P-HNN*, which achieved very strong results on a subset of three clinical datasets by modifying the Holistically-Nested Network (HNN) (Xie and Tu, 2015) structure to progressively sum side-output predictions during the decoder phase. In this study, we test the

efficacy of the proposed *SegCaps* algorithm for pathological lung segmentation due to precise segmentation's importance as a precursor to the deployment of nearly any computer-aided diagnosis (CAD) tool for pulmonary image analysis.

2.4. Muscle and Adipose (Fat) Tissue Segmentation from MRI

A number of applications favor MRI as the primary imaging modality, including most popularly cardiac applications. For example, Mortazi et al. (2017b) proposed a multi-view CNN, following an encoder-decoder structure and adding a novel loss function, for segmenting the left atrium and proximal pulmonary veins from MRI. Body composition analysis (e.g. segmenting/quantifying muscle and adipose (fat) tissue) favors MRI as well, due to its excellent soft tissue contrast and lack of ionizing radiation. In Irmakci et al. (2018), the authors proposed a method based on fuzzy connectivity to perform segmentation of muscle and fat tissue of the thigh region of whole-body MRI scans. This work represents the current state of the art results in terms of Dice score.

2.5. Capsule Networks

A simple three-layer capsule network, called *CapsNet*, showed remarkable initial results in Sabour et al. (2017), producing state-of-the-art classification results on the MNIST dataset and relatively good classification results on the CIFAR10 dataset. Since then, researchers have begun extending the idea of capsule networks to other applications, including brain-tumor classification (Afshar et al., 2018), lung-nodule screening (Mobiny and Van Nguyen, 2018), action detection (Duarte et al., 2018), point-cloud autoencoders (Zhao et al., 2019), adversarial detection (Frosst et al., 2018; Qin et al., 2019), and even creating wardrobes (Hsiao and Grauman, 2018), as well as several technical contributes to improve the routing mechanism for datasets such as MNIST, CIFAR10, SVHN, SmallNorb, etc. (Hinton et al., 2018; Kosiorek et al., 2019). Nonetheless, the majority of these works remain focused on small image classification, and no work yet exists in literature for a method of capsule-based object segmentation.

The remainder of the paper is organized as follows: Section 1.2 introduces the locally-constrained dynamic routing and transformation matrix sharing which are the key building blocks for our method; Section 3 describes our proposed *SegCaps* framework in detail, including the deconvolutional capsules and reconstruction regularization for segmentation. The network is a deep encoder-decoder architecture with skip connections concatenating together capsule types from earlier layer with the same spatial dimensions. The input and outputs shown are from the task of muscle segmentation from MRI scans of patient's thighs. segmentation; Section 4.3 details the five experimental datasets, our implementation settings (e.g. hyperparameters), and the results of our main experiments; Section 5 covers the ablation studies performed which help to determine the contribution of each aspect of our proposed method to the final results; and finally Section 6 is the discussion and conclusions of our work. Experimental results of our method applied to other types of imaging data and applications to provide empirical support for the general applicability of our study are included in the appendix.

3. SegCaps: Capsules for Object Segmentation

In the following section, we describe the formulation of our *SegCaps* architecture. As illustrated in Figure 2, the input to our *SegCaps* network is a large image (e.g. 512×512 pixels), in this case, a slice of a MRI Scan. The image is passed through a 2D convolutional layer which produces 16 feature maps of the same spatial dimensions. This output forms our first set of capsules, where we have a single capsule type with a grid of 512×512 capsules, each of which is a 16 dimensional vector. This is then followed by our first convolutional capsule layer. In the following, we generalize the process of our convolutional capsules and routing to any given layer ℓ in the network.

At layer ℓ , there exists a set of capsule types

$$T^\ell = \{t_1^\ell, t_2^\ell, \dots, t_n^\ell \mid n \in \mathbb{N}\}. \quad (1)$$

For every $t_i^\ell \in T^\ell$, there exists an $h^\ell \times w^\ell$ grid of z^ℓ -dimensional child capsules,

$$C = \{c_{1,1}, \dots, c_{1,w^\ell}, \dots, c_{h^\ell,1}, \dots, c_{h^\ell,w^\ell}\}, \quad (2)$$

where $h^\ell \times w^\ell$ is the spatial dimensions of the output of layer $\ell - 1$. At the next layer of the network, $\ell + 1$, there exists a set of capsule types

$$T^{\ell+1} = \{t_1^{\ell+1}, t_2^{\ell+1}, \dots, t_m^{\ell+1} \mid m \in \mathbb{N}\}. \quad (3)$$

And for every $t_j^{\ell+1} \in T^{\ell+1}$, there exists an $h^{\ell+1} \times w^{\ell+1}$ grid of $z^{\ell+1}$ -dimensional parent capsules,

$$P = \{p_{1,1}, \dots, p_{1,w^{\ell+1}}, \dots, p_{h^{\ell+1},1}, \dots, p_{h^{\ell+1},w^{\ell+1}}\}, \quad (4)$$

where $h^{\ell+1} \times w^{\ell+1}$ is the spatial dimensions of the output of layer ℓ .

In convolutional capsules, for every parent capsule type $t_j^{\ell+1} \in T^{\ell+1}$, every parent capsule $p_{x,y} \in P$ receives a set of ‘‘prediction vectors’’, $\{\hat{u}_{x,y|t_1^\ell}, \hat{u}_{x,y|t_2^\ell}, \dots, \hat{u}_{x,y|t_n^\ell}\}$, one for each capsule type in T^ℓ . This set of prediction vectors is defined as the matrix multiplication between a learned transformation matrix for the given parent capsule type, $M_{t_j^{\ell+1}}$, and the sub-grid of child capsules outputs, $U_{x,y|t_i^\ell}$, within a user-defined kernel centered at position (x,y) in layer ℓ ; hence

$$\hat{u}_{x,y|t_i^\ell} = M_{t_j^{\ell+1}} \cdot U_{x,y|t_i^\ell}, \quad \forall t_i^\ell \in T^\ell. \quad (5)$$

Explicitly, each $U_{x,y|t_i^\ell}$ has shape $k_h \times k_w \times z^\ell$, where $k_h \times k_w$ are the dimensions of the user-defined kernel, for all capsule types $t_i^\ell \in T^\ell$. Each $M_{t_j^{\ell+1}}$ has shape $k_h \times k_w \times z^\ell \times z^{\ell+1}$. Thus, we can see each $\hat{u}_{x,y|t_i^\ell}$ is a $z^{\ell+1}$ -dimensional vector, since these will be used to form our parent capsules. In practice, we solve for all parent capsule types simultaneously by defining M to have shape $k_h \times k_w \times z^\ell \times |T^{\ell+1}| \times z^{\ell+1}$, where $|T^{\ell+1}|$ is the number of parent capsule types in layer $\ell+1$. Note, as opposed to CapsNet, we are sharing transformation matrices across members of the grid (i.e. each $M_{t_j^{\ell+1}}$ does not depend on the spatial location (x,y)), as the same transformation matrix is shared across all spatial locations within a given capsule type, similar to how convolutional kernels scan an input feature map. This is one way our method can exploit parameter sharing to dramatically cut down on the total number of parameters to be learned. The values of these transformation matrices for each capsule type in a layer are learned via the backpropagation algorithm with a supervised loss function.

Algorithm 1 Locally-Constrained Dynamic Routing.

```

1: procedure ROUTING ( $\hat{u}_{x,y|t_i^\ell}, d, \ell, x, y$ )
2:   for all capsule types  $t_i^\ell$  at position  $(x,y)$  and capsule type  $t_j^{\ell+1}$  at position  $(x, y)$ :  $b_{t_i^\ell| x, y} \leftarrow 0$ .
3:   for  $d$  iterations do
4:     for all capsule types  $t_i^\ell$  at position  $(x,y)$ :  $\mathbf{r}_{t_i^\ell} \leftarrow \text{softmax}(\mathbf{b}_{t_i^\ell})$   $\triangleright$  softmax computes Eq. 7
5:     for all capsule types  $t_j^{\ell+1}$  at position  $(x, y)$ :  $\mathbf{p}_{x,y} \leftarrow \sum_n r_{t_i^\ell} | x, y \hat{u}_{x,y|t_i^\ell}$ 
6:     for all capsule types  $t_j^{\ell+1}$  at position  $(x, y)$ :  $\mathbf{v}_{x,y} \leftarrow \text{squash}(\mathbf{p}_{x,y})$   $\triangleright$  squash computes Eq. 8
7:     for all capsule types  $t_i^\ell$  and all capsule types  $t_j^{\ell+1}$ :  $b_{t_i^\ell| x, y} \leftarrow b_{t_i^\ell| x, y} + \hat{u}_{x,y|t_i^\ell} \cdot \mathbf{v}_{x,y}$ 
   return  $\mathbf{v}_{x,y}$ 

```

To determine the final input to each parent capsule $\mathbf{p}_{x,y} \in P$, where again P is the grid of parent capsules for parent capsule type $t_j^{\ell+1} \in T^{\ell+1}$, we compute the weighted sum over these “prediction vectors” as,

$$\mathbf{p}_{x,y} = \sum_n r_{t_i^\ell| x, y} \hat{u}_{x,y|t_i^\ell}, \quad (6)$$

where $r_{t_i^\ell| x, y}$ are the routing coefficients determined by the dynamic routing algorithm, and each member of the grid (x,y) has a unique routing coefficient. These routing coefficients are computed by a “routing softmax”,

$$r_{i_i^\ell | x, y} = \frac{\exp(b_{i_i^\ell | x, y})}{\sum_{t_j^\ell + 1} \exp(b_{i_i^\ell | t_j^\ell + 1})}, \quad (7)$$

whose initial logits, $b_{i_i^\ell | x, y}$ are the log prior probabilities that prediction vector $\hat{\mathbf{u}}_{x, y | t_i^\ell}$ should be routed to parent capsule $\mathbf{p}_{x, y}$. Note that the $\sum_{t_j^\ell + 1}$ term is across parent capsule types in $T^{\ell + 1}$ for each (x, y) location.

Our method differs from the dynamic routing implemented by Sabour et al. (2017) in two ways. First, we locally constrain the creation of the prediction vectors. Second, we only route the child capsules within the user-defined kernel to the parent, rather than routing every single child capsule to every single parent. The output capsule is then computed using a non-linear squashing function

$$\mathbf{v}_{x, y} = \frac{\|\mathbf{p}_{x, y}\|^2 \mathbf{p}_{x, y}}{1 + \|\mathbf{p}_{x, y}\|^2 \|\mathbf{p}_{x, y}\|}, \quad (8)$$

where $\mathbf{v}_{x, y}$ is the vector output of the capsule at spatial location (x, y) and $\mathbf{p}_{x, y}$ is its final input. Lastly, the agreement is measured as the scalar product,

$$a_{x, y | t_i^\ell} = \mathbf{v}_{x, y} \cdot \hat{\mathbf{u}}_{x, y | t_i^\ell}. \quad (9)$$

The pseudocode for this locally-constrained dynamic routing is summarized in Algorithm 1. A final segmentation mask is created by computing the length of the capsule vectors in the final layer and assigning the positive class to those whose magnitude is above a threshold, and the negative class otherwise.

3.1. Deconvolutional Capsules

In order to form a deep encoder-decoder network, we introduce the concept of “deconvolutional” capsules. These are similar to the locally-constrained convolutional capsules; however, the prediction vectors are now formed using the transpose of the operation previously described. Note that the dynamic routing of these differently-formed prediction vectors still occurs in the exact same way, so we will not re-describe that part of the operation.

The set of prediction vectors for deconvolutional capsules are defined again as the matrix multiplication between a learned transformation matrix, $M_{t_j^\ell + 1}$, for a given parent capsule type $t_j^\ell + 1 \in T^{\ell + 1}$, and the sub-grid of child capsules outputs, $W_{x, y | t_i^\ell}$ for each capsule type in $t_i^\ell \in T^\ell$, within a user-defined kernel centered at position (x, y) in layer ℓ . However, in deconvolutional capsules, we first need to reshape our child capsule outputs following the fractional striding formulation used in Long et al. (2015). This allows us to effectively upsample the height and width of our capsule grids by the scaling factor chosen. For each member of the grid, we can then form our prediction vectors again by

$$\widehat{\mathbf{w}}_{x,y|t_i^\ell} = \mathbf{M}_{t_j^\ell+1} \cdot \mathbf{W}_{x,y|t_i^\ell}, \quad \forall t_i^\ell \in T^\ell. \quad (10)$$

Thus, we have each $\widehat{\mathbf{w}}_{x,y|t_i^\ell}$ as a $z^{\ell+1}$ -dimensional vector, and is input to the dynamic routing algorithm to form our parent capsules. As before, in practice we solve for all parent capsule types simultaneously by defining M to have shape $k_h \times k_w \times z^\ell \times |T^{\ell+1}| \times z^{\ell+1}$, where $|T^{\ell+1}|$ is the number of parent capsule types in layer $\ell+1$. Here, we still sharing transformation matrices across members of the grid (i.e. each $M_{t_j^\ell+1}$ does not depend on the spatial location (x,y)), similar to how transposed convolutional kernels scan an input feature map.

3.2. Reconstruction Regularization

As a method of regularization, we extend the idea of reconstructing the input to promote a better embedding of our input space. This forces the network to not only retain all necessary information about a given input, but also encourages the network to better represent the full distribution of the input space, rather than focusing only on its most prominent modes relevant to the desired task. Since we only wish to model the distribution of the positive input class and treat all other pixels as background, we mask out segmentation capsules which do not belong to the positive class and reconstruct a similarly masked version of the input image. We perform this reconstruction via a three layer 1×1 convolutional network, then compute a mean-squared error (MSE) loss between only the positive input pixels and this reconstruction. More explicitly, we formulate this problem as

$$R^{x,y} = I^{x,y} \times S^{x,y} \mid S^{x,y} \in \{0, 1\}, \text{ and} \quad (11)$$

$$\mathcal{L}_R = \frac{\gamma}{X \times Y} \sum_x^X \sum_y^Y \|R^{x,y} - O_r^{x,y}\|, \quad (12)$$

where \mathcal{L}_R is the supervised loss for the reconstruction regularization, γ is a weighting coefficient for the reconstruction loss, $R^{x,y}$ is the reconstruction target pixel, $I^{x,y}$ is the image pixel, $S^{x,y}$ is the ground-truth segmentation mask value, and $O_r^{x,y}$ is the output of the reconstruction network, each at pixel location (x,y) , respectively, and X and Y are the width and height, respectively, of the input image. For simplicity γ was initially set to 1; however, the parameter produces similar results for settings from $1 - 0.001$. Lower than this or higher than this starts to degrade performance. An ablation study of the contribution of this regularization is included in Section 5. The total loss is the summation of this reconstruction loss and a weighted binary cross-entropy (BCE) loss for the segmentation output, weighted by the foreground/background pixel balance of each training set respectively.

4. Experiments and Results

4.1. Pathological Lung Datasets

Experiments were conducted on five pathological lung datasets, obtained from both clinical and pre-clinical subjects, containing nearly 2000 CT scans, with annotations by expert radiologists. An example typical scan with ground-truth from each dataset is shown in Figure 3. The three clinical and two pre-clinical (mice) datasets analyzed are as follows:

- *The Lung Image Database Consortium and Image Database Resource Initiative* (Armato et al., 2011), abbreviated as **LIDC-IDRI**, contains 885 annotated CT scans of lung cancer screening patients collected from seven academic centers and eight medical imaging companies. Scans were captured using seven different GE Medical Systems LightSpeed scanner models, four different Philips Brilliance scanner models, five different Siemens Denition, Emotion, and Sensation scanner models, and Toshiba Aquilion scanners. Slice thicknesses range from 0.6 to 5 mm with in-plane pixel sizes ranging from 0.461 to 0.977 mm. All image sizes are 512×512 pixels per slice.
- *The Lung Tissue Research Consortium database* (Karwoski et al., 2008), abbreviated as **LTRC**, contains 545 annotated CT scans, with most donor subjects having interstitial fibrotic lung disease or chronic obstructive pulmonary disease (COPD). All scans using the LTRC protocol were obtained using either General Electric or Siemens scanners with 16 or more detectors, and imaging parameters were standardized as much as possible among the enrollment centers (with slice thickness 1.25 mm or less with 50% overlapping reconstruction in a high-spatial, frequency-preserving algorithm). All image sizes are 512×512 pixels per slice.
- *The Multimedia Database of Interstitial Lung Diseases* (Depeursinge et al., 2012), abbreviated as **UHG**, built at the University Hospitals of Geneva contains 214 annotated CT scans of patients affected with one of the 13 histological diagnoses of interstitial lung disease (ILD). The dataset follows the HRCT scanning protocol, with a slice thickness of 1 – 2 mm, spacing between slices of 10 – 15 mm, scan time of 1 – 2 s, no contrast agent, axial pixel matrix of 512×512 , and x, y spacing of 0.4 – 1 mm.
- *The TB-Smoking dataset* collected at Johns Hopkins University, abbreviated as **JHU-TBS**, contains 108 annotated CT scans of mice subjects affected with tuberculosis (TB) and exposed to smoke inhalation. Slice thicknesses range from 0.1 to 0.2 mm with in-plane pixel sizes ranging from 0.1 to 0.2 mm. Images range in size from 176×176 to 352×352 pixels per slice.
- *The TB dataset* also collected at Johns Hopkins University, abbreviated as **JHU-TB**, contains 208 annotated CT scans of mice subjects affected with TB undergoing experimental treatment. Slice thicknesses range from 0.041 to 0.058 mm with in-plane pixel sizes ranging from 0.041 to 0.058 mm. Images range in size from 199×212 to 580×496 pixels per slice.

In total, 1960 CT scans were annotated in this study. Each dataset was treated completely separate, as each offers unique challenges to automated segmentation algorithms. For preprocessing, all CT scans were clipped at -1024 and 3072 , then normalized to a $0 - 1$ scale. All images used their original resolutions during training and testing. Ten-fold cross-validation was performed for training all algorithms, with 10% of training data left aside for validation and early-stopping. The mean and standard deviation (std) across the 10-folds for each dataset is presented for two key metrics, namely the 3D Dice similarity coefficient (Dice) and 3D Hausdorff distance (HD) computer for each 3D CT scan.

4.2. Implementation Details

All algorithms, namely *U-Net*, *Tiramisu*, *P-HNN*, our three-layer baseline capsule segmentation network, and *SegCaps* are all implemented using Keras (Chollet et al., 2015) with TensorFlow (Abadi et al., 2015). The *U-Net* architecture is implemented exactly as described in the original paper by Ronneberger et al. (2015). *P-HNN* was implemented based on their official Caffe code, including individual layer-specific learning rate multipliers and kernel initialization. However, we removed the layer-specific learning rate and changed the kernel initialization to Xavier to match the other networks and achieve much better results. *Tiramisu* follows the highest performing model presented in (Jégou et al., 2017), namely *FC-DenseNet103*. To remain consistent, since pre-trained models are not available for our custom-designed *SegCaps*, and to better see the performance of each individual method under different amounts of training data and pathologies present, no pre-trained weights were used to initialize any of the models; instead, all were trained from scratch on each dataset investigated. It can be reasonably assumed based on previous studies that pre-training on large datasets such as ImageNet would improve the performance of all models. A weighted-BCE loss is used for the segmentation output of all networks, with weights determined by the foreground/background pixel balance of each training set respectively. For the capsule network, the reconstruction output loss is computed via the masked-MSE described in Section 3. All possible experimental factors are controlled between different networks; all networks are trained from scratch, using the same data augmentation methods (scale, flip, shift, rotate, elastic deformations, and random noise) and Adam optimization (Kingma and Ba, 2014) with an initial learning rate of 0.00001. A batch size of 1 is chosen for all experiments to match the original *U-Net* implementation. The learning rate is decayed by a factor of 0.05 upon validation loss stagnation for 50,000 iterations and early-stopping is performed with a patience of 250,000 iterations based on validation 2D Dice scores. Positive/negative pixels were set in the segmentation masks based on a threshold of on the networks' output score maps. Thresholds are found dynamically for each testing split based on which level provides the best Dice score for the validation set of images. All code is made publicly available.¹

4.3. Lung Segmentation Results

The final quantitative results of these experiments to perform lung segmentation from pathological CT scans are shown in Tables 1 - 5. Table 1 shows results on the LIDC-IDRI dataset, the largest of the three clinical datasets with typically the least severe pathology

¹<https://github.com/lalonderodney/SegCaps>

present on average compared to the other two clinical datasets. Table 2 shows results on the LTRC dataset, a large dataset with large amounts of ILD and COPD pathology present. Table 3 shows results on the UHG dataset, perhaps the most challenging of the three clinical datasets, both due to its relatively smaller size and the severe average amount of pathology present in patients scanned. Table 4 shows results on the JHU-TBS dataset, and provides the first fully-automated deep learning based segmentation results presented in the literature for lung segmentation on pre-clinical subjects. Table 5 shows results on the JHU-TB dataset, a larger but more challenging dataset of mouse subjects with typically more severe pathology present than the JHU-TBS dataset.

The results of these experiments show *SegCaps* consistently outperforms all other compared state-of-the-art approaches in terms of the commonly measured metrics, Dice and HD. Additionally, *SegCaps* achieves this while only using a fraction of the total parameters of these much larger networks. The proposed *SegCaps* architecture contains less than 4.6% parameters than *U-Net*, less than 9.5 % of *P-HNN*, and less than 14.9 % of *Tiramisu*. A comparison with similarly sized version of these other networks is shown in Section 5.2. As a brief note in regardless to the discrepancy in results for *P-HNN* between our study and those in the original work, this can be explained by several factors: the original work i) used ImageNet pre-trained models, ii) selected a carefully chosen subset (73 scans) of the UHG dataset, and iii) trained and tested models using all datasets combined in the cross-validation splits.

Qualitative results for typical samples from all datasets are shown in Figures 4 – 8. As can be seen in these qualitative examples, *SegCaps* achieves higher results by not falling into the typical segmentation failure causes, namely over-segmentation and segmentation-leakage. These qualitative examples are supported by our quantitative findings where over-segmentation is best captured by the HD metric and segmentation-leakages are best captured by the Dice metric.

Further, we investigate how different capsule vectors in the final segmentation capsule layer are representing different visual attributes. Figure 9 shows three selected visual attributes (each row) out of the sixteen (dimension of final capsule segmentation vector) across different perturbation values of the vectors ranging from -0.25 to $+0.25$ (each column) for an example clinical and pre-clinical scan. We observe that regions with different textural properties (i.e., small and large homogeneous) are progressively captured by the different dimensions of the capsule segmentation vectors.

4.4. Muscle and Adipose (Fat) Tissue Segmentation Datasets and Preprocessing

Experiments were conducted on the Baltimore Longitudinal Study of Ageing (BLSA) (Ferrucci, 2008), where a total of 150 scans were collected using three contrasts from 50 subjects. These MRI were acquired using a 3T Philips Achieva MRI scanner (Philips Healthcare, Best, The Netherlands) equipped with a Q-body radiofrequency coil for transmission and reception. Three different T1-weighted MR contrasts, namely water and fat, water-only (fat-suppressed), and fat-only (water-suppressed), were used, where water and fat suppression were achieved using spectral pre-saturation with inversion recovery (SPIR), with coverage from the proximal to distal ends of the femur using 80 slices in the

foot to head direction, a field of view (FOV) of $440 \times 296 \times 400 \text{ mm}^3$ and a voxel size of $1 \times 1 \text{ mm}^2$ in-plane, and slice thickness varies from 1 mm to 3 mm in different scans (one particular scan was with 5 mm slice thickness). The age of subjects ranged between 44 – 89 years and the body mass index (BMI) ranged from 18.67 – 45.68. Examples of each MRI contrast with the ground truth-annotations (GT) for both muscle and adipose (fat) tissue are shown in Figure 10.

For training and testing, we performed preprocessing on the MRI images. First we applied the non-uniform non-parametric intensity normalisation technique by Tustison et al. (2010) to remove field bias. Next we apply Curvature Anisotropic Diffusion to smooth the image and remove noise. Lastly we perform z-score normalization to normalize the intensities before standardizing the image to the 0 – 1 range.

4.5. Muscle and Adipose (Fat) Tissue Segmentation Results

Experiments were performed using U-Net and SegCaps and compared to the state-of-the-art method by Irmakci et al. (2018) using the same comparative metric (*i.e.* Dice coefficient). The results of these experiments are shown in Tables 6 – 8 and show that both U-Net and SegCaps can outperform the previous state-of-the-art, while SegCaps does so using again only a small fraction of the parameters as U-Net while performing at the same level as U-Net. The results reported for Irmakci et al. (2018) are the results from the original work (using manual seeding) on the same dataset. For U-Net and SegCaps, qualitative results are shown in Fig 11 for six of the 50 patients in the BLSA dataset. As with the lung segmentation experiments, U-Net tends to struggle with areas of similar intensity values, which do not belong to the correct tissue class.

4.6. Generalizing to Unseen Orientations of Objects

In a final set of experiments, we tested the affine equivariant property of capsule networks on natural images. It has been stated that, due to the affine projections of capsule vectors from children to parents, capsules should be robust to affine transformations on the input, and should in fact be able to generalize to *unseen* poses of target classes. However, no study has formally demonstrated this property. In this experiment, we randomly selected images from the PASCAL VOC dataset which contained only a single foreground object. Both *U-Net* and *SegCaps* were then trained on a single selected image until training accuracy converge to 100%, which occurred around 1000 epochs for both networks. For training, each network followed exactly the training settings described in Section 4.2. Each network was then tested on 90 degree rotations and the mirroring of the training image. *SegCaps* performed well on nearly all images tested, while *U-Net* performed quite poorly, as can be seen in Figure 12. Since *U-Net* has significantly more parameters than *SegCaps*, we also ran experiments at 10000 epochs, long after both networks had converged to 100% training accuracy. *U-Net* continued to present failures, where *SegCaps* did not suffer the same issue. This shows that *SegCaps* is indeed far more robust to affine transformations on the input, a significant issue for CNNs as shown in both this experiment and works such as by Alcorn et al. (2019).

5. Ablations Studies

In the following subsections, we investigate the role of the deeper encoder-decoder network structure enabled by the introduction of our deconvolutional capsules, the effect of the reconstruction regularization, the optimal number of dynamic routing iterations to perform, and the relative efficiency of parameter use with similarly-sized versions of all studied networks. The UHG dataset is perhaps the most challenging of the three clinical lung segmentation datasets in our study, both due to its relatively smaller size and the average amount of pathology present in patients scanned. As seen in Table 3, results on all metrics are significantly lower for this challenging dataset. For those reasons, and the lower performance scores leading to bigger differences between approaches, as well as the dataset being publicly available, we chose this dataset for running our ablation experiments.

5.1. Network Structure/Deconvolutional Capsules

The original CapsNet introduced by Sabour et al. (2017) was a simple three layer network, consisting of a single convolutional layer, a primary capsule layer (convolutional layer with a reshape function), and a fully-connected capsule layer. This network achieved remarkable results for its size, beating the state-of-the-art on MNIST and performing well on CIFAR10. In our initial efforts for this study, we attempted to apply this network to the task of segmentation, however, the fully-connected capsule layer was far too memory intensive to make this approach viable with our 512×512 2D slices of CT scans. After introducing the locally-constrained dynamic routing and transformation matrix sharing, we then created a network nearly identical to the original CapsNet with the fully-connected capsule layer swapped out for our locally-constrained version. A diagram of this network is shown in Figure 1. The results of this network on the UHG dataset is shown in Table 9. As one might expect, swapping out a layer which is fully-connected in space for one which is locally-connected dramatically hurt the performance for a task which relies on global information (*i.e.* determining lung tissue/air from non-lung tissue, bone, etc.). This motivated the introduction of the “deconvolutional” capsule layer which allows for the creation of deep encoder-decoder networks, and thus the recovery of global information, retention of local information, and the parameter savings of locally-constrained capsules.

5.2. Parameter Use

Shown in Tables 10– 11, we investigate the number of parameters in the proposed *SegCaps*, *U-Net*, *Tiramisu* and *P-HNN*, as well as down-scaled versions of *U-Net*, *Tiramisu*, and *P-HNN*. *U-Net* and *P-HNN* are scaled down by dividing the number of feature maps per layer by a constant factor, $k = 4.68$ and $k = 3.2$ respectively, and *Tiramisu* is scaled down by using the lighter FC-DenseNet56 purposed in the original work by Jégou et al. (2017). When the parameters of *U-Net* and *P-HNN* are scaled down to roughly the same number of parameters as *SegCaps*, these models perform comparatively worse, as shown in Table 11, providing evidence that *SegCaps* is able to make better use of the parameters available to it than its CNN counterparts. *Tiramisu-56* is a minor exception to this trend as its Dice score remained similar while the HD only fell slightly from *Tiramisu-103*. The reason for this is most likely because *Tiramisu-56* was carefully engineered to achieve the highest possible accuracy with few parameters while the addition of dense connections has been shown to make far better

use of parameters than standard non-dense CNNs (Huang et al., 2017). However, as can be seen in Table 11, when all networks have roughly the same number of parameters, *SegCaps* outperforms all other methods.

5.3. Reconstruction Regularization

The idea of reconstructing the input as a method of regularization was used in CapsNet by Sabour et al. (2017). The theory behind this technique and the regularization effect it introduces is similar in nature to the problem of “mode collapse” in generative adversarial networks (GANs). When training a generative neural network for a specific task through the backpropagation algorithm, the model “collapses” to focusing on only the most prevalent modes in the data distribution. A similar phenomenon occurs when you train a discriminative network for a specific task, the model “collapses” to only focus on the most discriminative features in the input data and ignores all others. By mapping the capsule vectors back to the input data, this forces the network to pay attention to more relevant features about the input, which might not be as discriminative for the given task, yet still provide some useful information, as evident by the improved results shown in Table 12. A similar results can be seen in VEEGAN by Srivastava et al. (2017), where they help solve the issue of mode collapse in GANs through a reconstructor network which reverses the action of the generator by mapping from data to noise.

5.4. Dynamic Routing Iterations

Since the dynamic routing algorithm chosen for this study is an iterative process, we can investigate the optimal number of times to run the routing algorithm per forward pass of the network. In the original work by Sabour et al. (2017), they found three iterations to provide the optimal results. As seen in Table 13, the number of routing iterations does have an effect on the network’s performance, and we find the same result in this study of three iterations being optimal over a set of different numbers of iterations studied. Several other recent studies have also found three routing iterations to achieve optimal performance (LaLonde et al., 2020a,b). However, other recent studies have found different routing methods or number of routing iterations to be optimal. Likely, as found by Paik et al. (2019), capsule networks will likely need to find an improved routing mechanism.

6. Discussions & Conclusion

We propose a novel deep learning algorithm, called *SegCaps*, for biomedical image segmentation, and showed its efficacy in a challenging problem of pathological lung segmentation from CT scans and thigh muscle and adipose (fat) tissue segmentation from MRI scans, as well as experiments around the affine equivariance properties of a capsule-based segmentation network. The proposed framework is the first use of the recently introduced capsule network architecture and expands it in several significant ways. First, we modify the original dynamic routing algorithm to act locally when routing children capsules to parent capsules and to share transformation matrices across capsules within the same capsule type. These changes dramatically reduce the memory and parameter burden of the original capsule implementation and allows for operating on large image sizes, whereas previous capsule networks were restricted to very small inputs. To compensate for the loss of

global information, we introduce the concept of “deconvolutional capsules” and a deep convolutional-deconvolutional capsule architecture for pixel level predictions of object labels. Finally, we extend the masked reconstruction of the target class as a regularization strategy for the segmentation problem.

Experimentally, *SegCaps* produces improved accuracy for lung segmentation on five datasets from clinical and pre-clinical subjects, in terms of Dice coefficient and Hausdorff distance, when compared with state-of-the-art networks *U-Net* (Ronneberger et al., 2015), *Tiramisu* (Jégou et al., 2017), and *P-HNN* (Harrison et al., 2017). For muscle and adipose (fat) tissue segmentation, *SegCaps* can perform on par with *U-Net* while only using a small fraction of the parameters, and outperforms the previous state-of-the-art. More importantly, the proposed *SegCaps* architecture provides strong evidence that the capsule-based framework can more efficiently utilize network parameters, achieving higher predictive performance while using 95.4% fewer parameters than *U-Net*, 90.5% fewer than *P-HNN*, and 85.1% fewer than *Tiramisu*. To the best of our knowledge, this work represents the largest study in pathological lung segmentation, and the only showing results on pre-clinical subjects utilizing state-of-the-art deep learning methods.

The results of these experiments demonstrate the effectiveness of the proposed capsule-based segmentation framework. This study provides helpful insights into future capsule-based works and provides lung-field segmentation analysis on pre-clinical subjects for the first time in the literature.

Acknowledgments

This study is partially supported by the NIH grant R01-EB020539 and R01-CA246704.

References

- Abadi M, et al., 2015. TensorFlow: Large-scale machine learning on heterogeneous systems. Software available from [tensorflow.org](https://www.tensorflow.org).
- Afshar P, Mohammadi A, Plataniotis KN, 2018. Brain tumor type classification via capsule networks, in: 2018 25th IEEE International Conference on Image Processing (ICIP), pp. 3129–3133.
- Alcorn MA, et al., 2019. Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4845–4854.
- Armato SG, et al., 2011. The lung image database consortium (lidc) and image database resource initiative (idri): A completed reference database of lung nodules on ct scans. *Medical Physics* 38, 915–931. [PubMed: 21452728]
- Badrinarayanan V, Kendall A, Cipolla R, 2017. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence* 39, 2481–2495. [PubMed: 28060704]
- Bagci U, Chen X, Udupa JK, 2012. Hierarchical scale-based multiobject recognition of 3-d anatomical structures. *IEEE Transactions on Medical Imaging* 31, 777–789. [PubMed: 22203704]
- Chen LC, Papandreou G, Kokkinos I, Murphy K, Yuille AL, 2018a. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence* 40, 834–848. [PubMed: 28463186]
- Chen LC, Zhu Y, Papandreou G, Schroff F, Adam H, 2018b. Encoder-decoder with atrous separable convolution for semantic image segmentation, in: Proceedings of the European conference on computer vision (ECCV), pp. 801–818.

- Chollet F, et al., 2015. Keras. <https://keras.io>.
- Çiçek Ö, Abdulkadir A, Lienkamp SS, Brox T, Ronneberger O, 2016. 3d u-net: learning dense volumetric segmentation from sparse annotation, in: International conference on medical image computing and computer-assisted intervention, Springer. pp. 424–432.
- Elnakib A, Gimel'farb G, Suri JS, El-Baz A, 2011. Medical Image Segmentation: A Brief Survey. Springer New York, New York, NY. pp. 1–39.
- Farag AA, Ahmed MN, El-Baz A, Hassan H, 2005. Advanced Segmentation Techniques. Springer US, Boston, MA. pp. 479–533.
- Ferrucci L, 2008. The baltimore longitudinal study of aging (blsa): a 50-year-long journey and plans for the future.
- Irmakci I, Hussein S, Savran A, Kalyani RR, Reiter D, Chia CW, Fishbein KW, Spencer RG, Ferrucci L, Bagci U, 2018. A novel extension to fuzzy connectivity for body composition analysis: Applications in thigh, brain, and whole body tissue segmentation. *IEEE Transactions on Biomedical Engineering* 66, 1069–1081. [PubMed: 30176577]
- LaLonde R, Bagci U, 2018. Capsules for object segmentation. arXiv preprint arXiv:1804.04241.
- LaLonde R, Kandel P, Spampinato C, Wallace MB, Bagci U, 2020a. Diagnosing colorectal polyps in the wild with capsule networks, in: 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI), IEEE. pp. 1086–1090.
- LaLonde R, Torigian D, Bagci U, 2020b. Encoding visual attributes in capsules for explainable medical diagnoses, in: Medical Image Computing and Computer Assisted Intervention MICCAI, Springer International Publishing, Cham.
- Milletari F, Navab N, Ahmadi SA, 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation, in: 2016 fourth international conference on 3D vision (3DV), IEEE. pp. 565–571.
- Paik I, Kwak T, Kim I, 2019. Capsule networks need an improved routing algorithm. arXiv preprint arXiv:1907.13327.
- Tustison NJ, Avants BB, Cook PA, Zheng Y, Egan A, Yushkevich PA, Gee JC, 2010. N4itk: improved n3 bias correction. *IEEE transactions on medical imaging* 29, 1310–1320. [PubMed: 20378467]
- Xu J, Xue K, Zhang K, 2019. Current status and future trends of clinical diagnoses via image-based deep learning. *Theranostics* 9, 7556. [PubMed: 31695786]
- Dalal N, Triggs B, 2005. Histograms of oriented gradients for human detection, in: Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on, IEEE. pp. 886–893.
- Depeursinge A, Vargas A, Platon A, Geissbuhler A, Poletti PA, Müller H, 2012. Building a reference multimedia database for interstitial lung diseases. *Computerized Medical Imaging and Graphics* 36, 227 – 238. [PubMed: 21803548]
- Duarte K, Rawat Y, Shah M, 2018. Videocapsulenet: A simplified network for action detection, in: Advances in Neural Information Processing Systems, pp. 7610–7619.
- El-Baz A, Farag A, Gimel'farb G, Falk R, El-Ghar MA, Eldiasty T, 2006. A framework for automatic segmentation of lung nodules from low dose chest ct scans, in: 18th International Conference on Pattern Recognition (ICPR'06), pp. 611–614.
- Felzenszwalb PF, Huttenlocher DP, 2004. Efficient graph-based image segmentation. *International journal of computer vision* 59, 167–181.
- Frosst N, Sabour S, Hinton G, 2018. Darccc: Detecting adversaries by reconstruction from class conditional capsules. arXiv preprint arXiv:1811.06969.
- Grady L, 2006. Random walks for image segmentation. *IEEE transactions on pattern analysis and machine intelligence* 28, 1768–1783. [PubMed: 17063682]
- Harrison AP, Xu Z, George K, Lu L, Summers RM, Mollura DJ, 2017. Progressive and multi-path holistically nested neural networks for pathological lung segmentation from ct images, in: Medical Image Computing and Computer Assisted Intervention MICCAI 2017, Springer International Publishing, Cham. pp. 621–629.
- He K, Gkioxari G, Dollár P, Girshick R, 2017. Mask r-cnn, in: Proceedings of the IEEE international conference on computer vision, pp. 2961–2969.

- He K, Zhang X, Ren S, Sun J, 2016. Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778.
- Hinton GE, Sabour S, Frosst N, 2018. Matrix capsules with em routing, in: 6th International Conference on Learning Representations, ICLR.
- Horowitz J, Pavlidis T, 1974. Picture segmentation by a direct split-and-merge procedure, in: Proc. of the 2nd Int. Joint Conf. on Pattern Recognition, pp. 424–433.
- Hsiao WL, Grauman K, 2018. Creating capsule wardrobes from fashion images, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7161–7170.
- Hu J, Shen L, Sun G, 2018. Squeeze-and-excitation networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 7132–7141.
- Huang G, Liu Z, Weinberger KQ, van der Maaten L, 2017. Densely connected convolutional networks, in: Computer Vision and Pattern Recognition, 2017. CVPR 2017. IEEE Computer Society Conference on.
- Jégou S, Drozdal M, Vazquez D, Romero A, Bengio Y, 2017. The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation, in: IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on, IEEE. pp. 1175–1183.
- Karwoski RA, Bartholmai B, Zavaletta VA, Holmes D, Robb RA, 2008. Processing of ct images for analysis of diffuse lung disease in the lung tissue research consortium, in: Medical Imaging 2008: Physiology, Function, and Structure from Medical Images, International Society for Optics and Photonics. p. 691614.
- Kingma DP, Ba J, 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
- Kosiorok AR, Sabour S, Teh YW, Hinton GE, 2019. Stacked capsule autoencoders. arXiv preprint arXiv:1906.06818.
- LaLonde R, Zhang D, Shah M, 2018. Clusternet: Detecting small objects in large scenes by exploiting spatio-temporal information, in: Computer Vision and Pattern Recognition, 2018. CVPR 2018. IEEE Computer Society Conference on.
- Lin G, Milan A, Shen C, Reid I, 2017. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation, in: Computer Vision and Pattern Recognition, 2017. CVPR 2017. IEEE Computer Society Conference on.
- Long J, Shelhamer E, Darrell T, 2015. Fully convolutional networks for semantic segmentation, in: Computer Vision and Pattern Recognition, 2015. CVPR 2015. IEEE Computer Society Conference on, pp. 3431–3440.
- Lowe DG, 1999. Object recognition from local scale-invariant features, in: Computer vision, 1999. The proceedings of the seventh IEEE international conference on, Ieee. pp. 1150–1157.
- Mansoor A, et al., 2014. A generic approach to pathological lung segmentation. IEEE transactions on medical imaging 33, 2293–2310. [PubMed: 25020069]
- Mobiny A, Van Nguyen H, 2018. Fast capsnet for lung cancer screening, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 741–749.
- Mortazi A, Burt J, Bagci U, 2017a. Multi-planar deep segmentation networks for cardiac substructures from mri and ct. stat 1050, 3.
- Mortazi A, Karim R, Rhode K, Burt J, Bagci U, 2017b. Cardiacnet: Segmentation of left atrium and proximal pulmonary veins from mri using multi-view cnn, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 377–385.
- Peng C, Zhang X, Yu G, Luo G, Sun J, 2017. Large kernel matters—improve semantic segmentation by global convolutional network, in: Computer Vision and Pattern Recognition, 2017. CVPR 2017. IEEE Computer Society Conference on, pp. 4353–4361.
- Pham DL, Xu C, Prince JL, 2000. Current methods in medical image segmentation. Annual review of biomedical engineering 2, 315–337.
- Qin Y, Frosst N, Sabour S, Raffel C, Cottrell G, Hinton G, 2019. Detecting and diagnosing adversarial images with class-conditional capsule reconstructions. arXiv preprint arXiv:1907.02957.
- Ren S, He K, Girshick R, Sun J, 2015. Faster r-cnn: Towards real-time object detection with region proposal networks, in: Advances in neural information processing systems, pp. 91–99.

- Ronneberger O, Fischer P, Brox T, 2015. U-net: Convolutional networks for biomedical image segmentation, in: International Conference on Medical image computing and computer-assisted intervention, Springer. pp. 234–241.
- Rosenfeld A, Kak AC, 1982. Digital Picture Processing: Volume 1. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. 2 edition.
- Sabour S, Frosst N, Hinton GE, 2017. Dynamic routing between capsules, in: Advances in Neural Information Processing Systems, pp. 3859–3869.
- Srivastava A, Valkov L, Russell C, Gutmann MU, Sutton C, 2017. Veegan: Reducing mode collapse in gans using implicit variational learning, in: Advances in Neural Information Processing Systems, pp. 3308–3318.
- Udupa JK, Samarasekera S, 1996. Fuzzy connectedness and object definition: theory, algorithms, and applications in image segmentation. Graphical models and image processing 58, 246–261.
- Vese LA, Chan TF, 2002. A multiphase level set framework for image segmentation using the mumford and shah model. International journal of computer vision 50, 271–293.
- Xie S, Tu Z, 2015. Holistically-nested edge detection, in: Proceedings of IEEE International Conference on Computer Vision.
- Yang M, Yu K, Zhang C, Li Z, Yang K, 2018. Denseaspp for semantic segmentation in street scenes, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3684–3692.
- Zhao H, Shi J, Qi X, Wang X, Jia J, 2017. Pyramid scene parsing network, in: Computer Vision and Pattern Recognition, 2017. CVPR 2017. IEEE Computer Society Conference on, pp. 2881–2890.
- Zhao Y, Birdal T, Deng H, Tombari F, 2019. 3d point capsule networks, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1009–1018.

- First ever capsule network for image segmentation.
- Reduced memory burden: locally-constrained routing and transformation matrix sharing.
- Introduced novel deconvolutional capsules to create encoder-decoder architecture.
- Extended the reconstruction regularization to the segmentation task.
- Outperformed existing methods across five clinical and preclinical datasets for lung segmentation from CT scans.

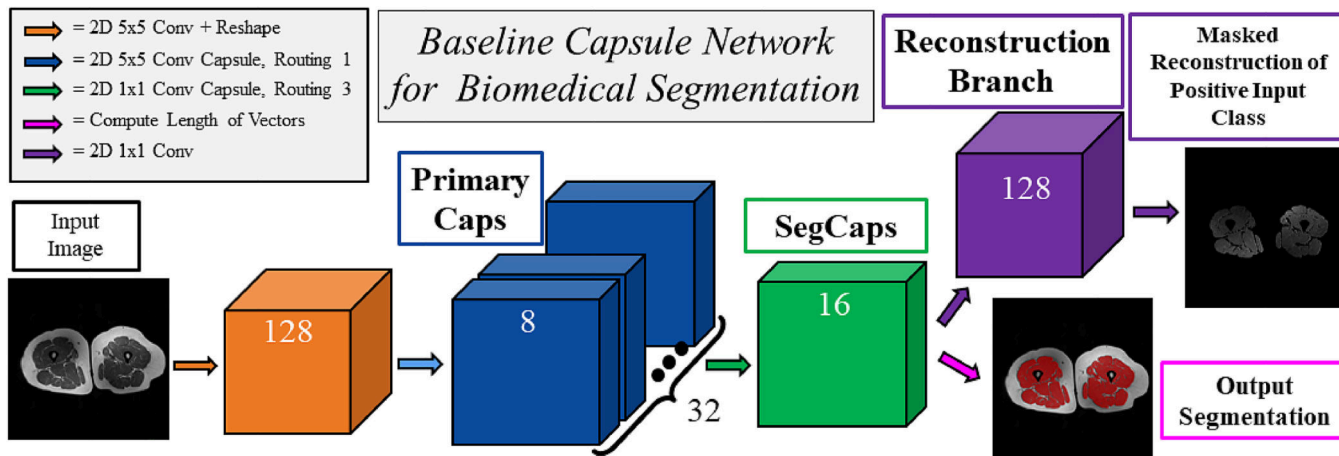


Figure 1: A simple three-layer capsule segmentation network closely mimicking the work by Sabour et al. (2017). This baseline capsule network uses our proposed locally-constrained dynamic routing algorithm with transformation matrix sharing, as well as the masked reconstruction of the positive input class. The input and outputs shown are of muscle tissue segmentation from MRI scans.

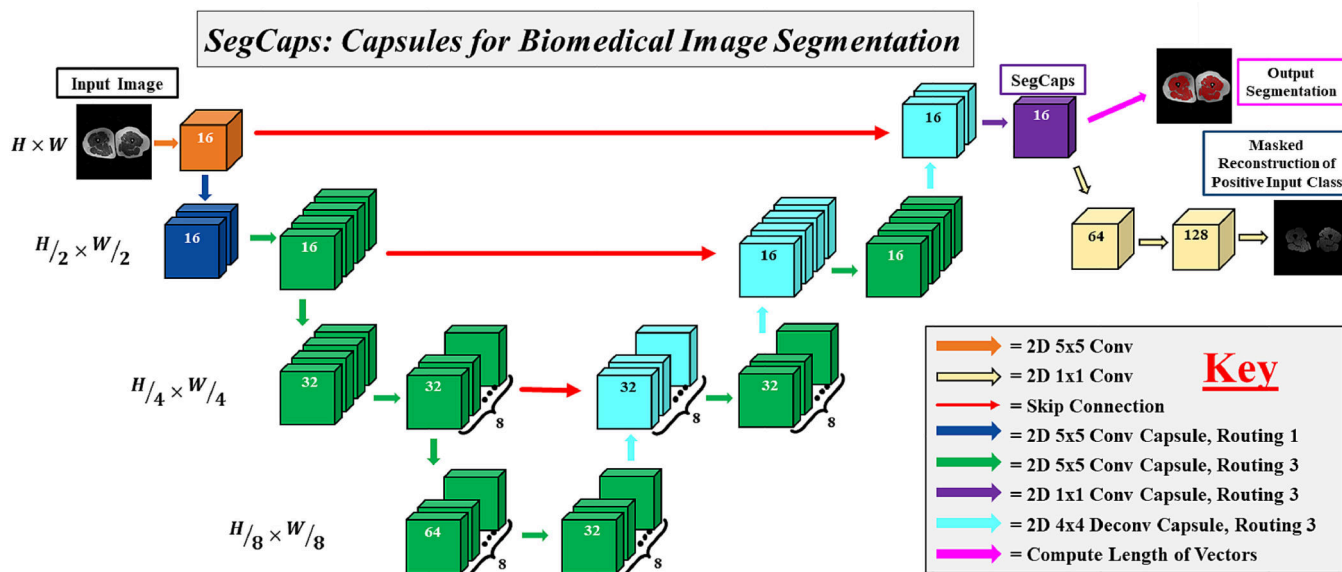


Figure 2:
The proposed *SegCaps* architecture for biomedical image




Figure 3:
Example scans with ground-truth masks (magenta) for each of the five datasets in this study.




Figure 4: Qualitative results for a 2D slice from a CT scan taken from the LIDC-IDRI dataset. It can be noticed that the CNN-based methods' typical failure cases are where the pixel intensities (Hounsfield units) are far from the class mean (*i.e.* high values within the lung regions or low values outside the lung regions).




Figure 5: Qualitative results for a 2D slice from a CT scan taken from the LTRC dataset. It can be noticed that the CNN-based methods' typical failure cases are where the pixel intensities (Hounsfield units) are far from the class mean (*i.e.* high values within the lung regions or low values outside the lung regions).




Figure 6: Qualitative results for a 2D slice from a CT scan taken from the UHG dataset. It can be noticed that the CNN-based methods' typical failure cases are where the pixel intensities (Hounsfield units) are far from the class mean (*i.e.* high values within the lung regions or low values outside the lung regions).




Figure 7: Qualitative results for a 2D slice from a CT scan taken from the JHU-TBS dataset. Note this drastically different anatomy and high level of noise present in the preclinical mice subjects. It can be noticed that the CNN-based methods' typical failure cases are where the pixel intensities (Hounsfield units) are far from the class mean (*i.e.* high values within the lung regions or low values outside the lung regions).




Figure 8: Qualitative results for a 2D slice from a CT scan taken from the JHU-TB dataset. Note this drastically different anatomy and high level of noise present in the preclinical mice subjects. It can be noticed that the CNN-based methods' typical failure cases are where the pixel intensities (Hounsfield units) are far from the class mean (*i.e.* high values within the lung regions or low values outside the lung regions).

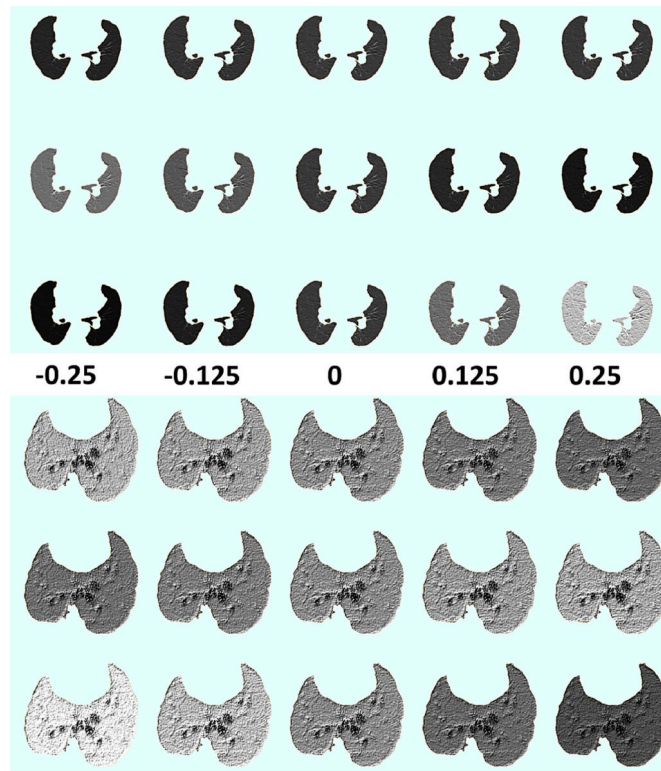


Figure 9: Reconstructions of selected capsule vectors (rows) under different perturbations from -0.25 to 0.25 (columns). The top three rows are reconstructions of a scan slice from the clinical LTRC dataset, while the bottom three are from the pre-clinical JHU-TB dataset. These results demonstrate that different dimensions of the capsule vectors are in fact learning different attributes of the lung tissue being segmented.

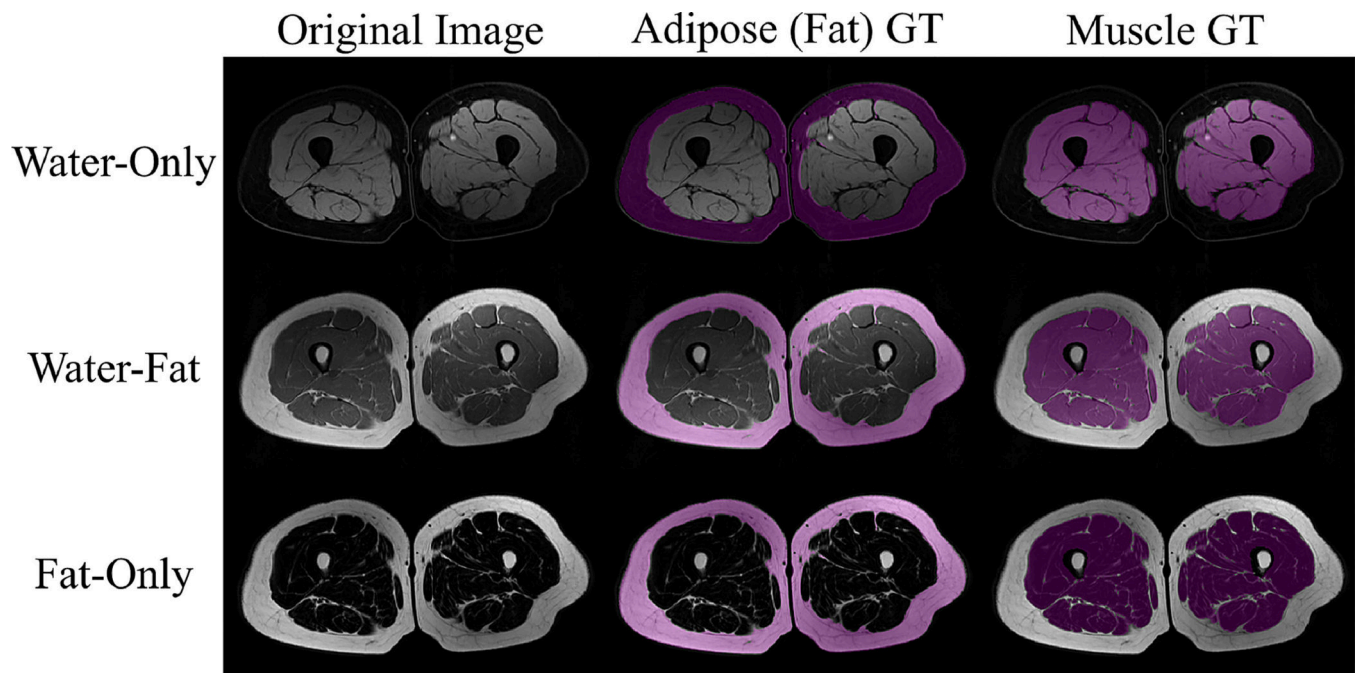


Figure 10: Example magnetic resonance (MR) images from the Baltimore Longitudinal Study of Ageing (BLSA) (Ferrucci, 2008) dataset. Three different T1-weighted MR contrasts, namely water and fat, water-only (fat-suppressed), and fat-only (water-suppressed) are shown with their ground-truth (GT) annotations.

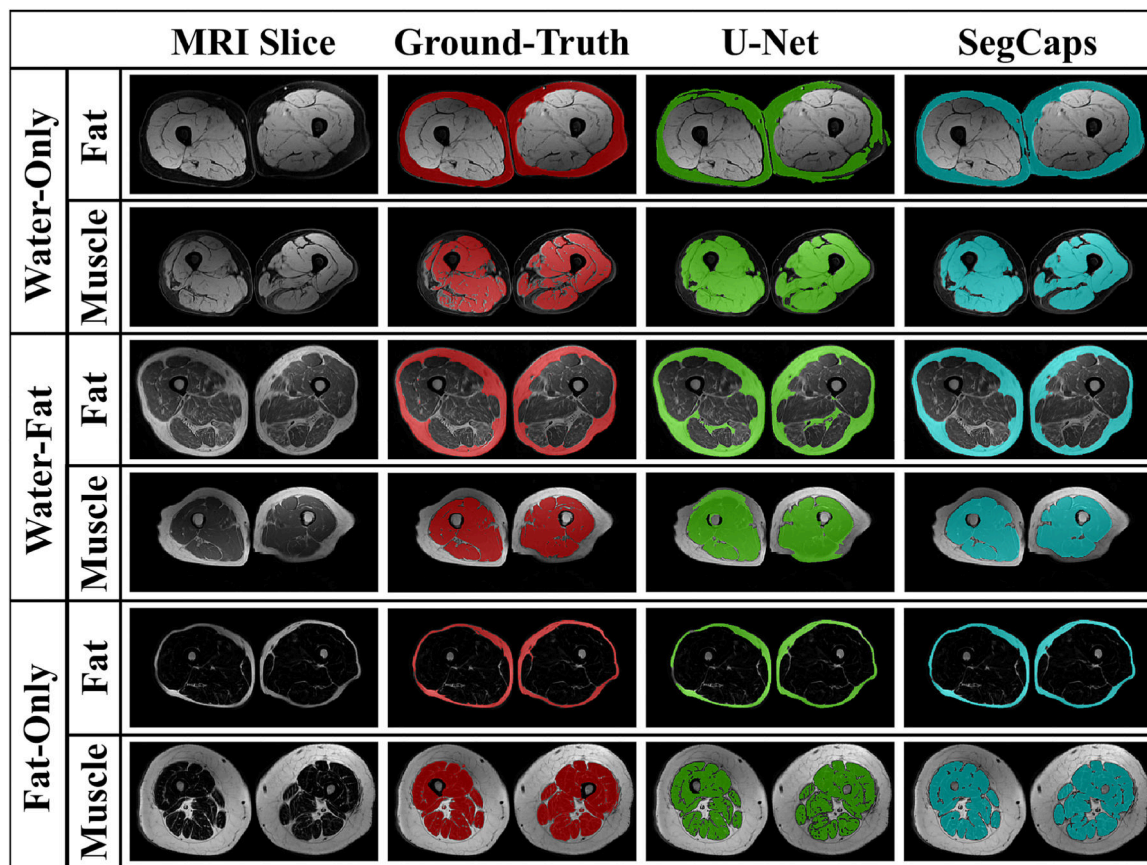


Figure 11:

Qualitative results on the BLSA dataset for T1-weighted water-only (fat-suppressed), water-fat, and fat-only (water-suppressed) contrasts of six different patients. Results are shown for U-Net and SegCaps, with U-Net showing systematic issues with areas of similar intensities values to foreground class, but which actually belong to the background class, and vice-versa.

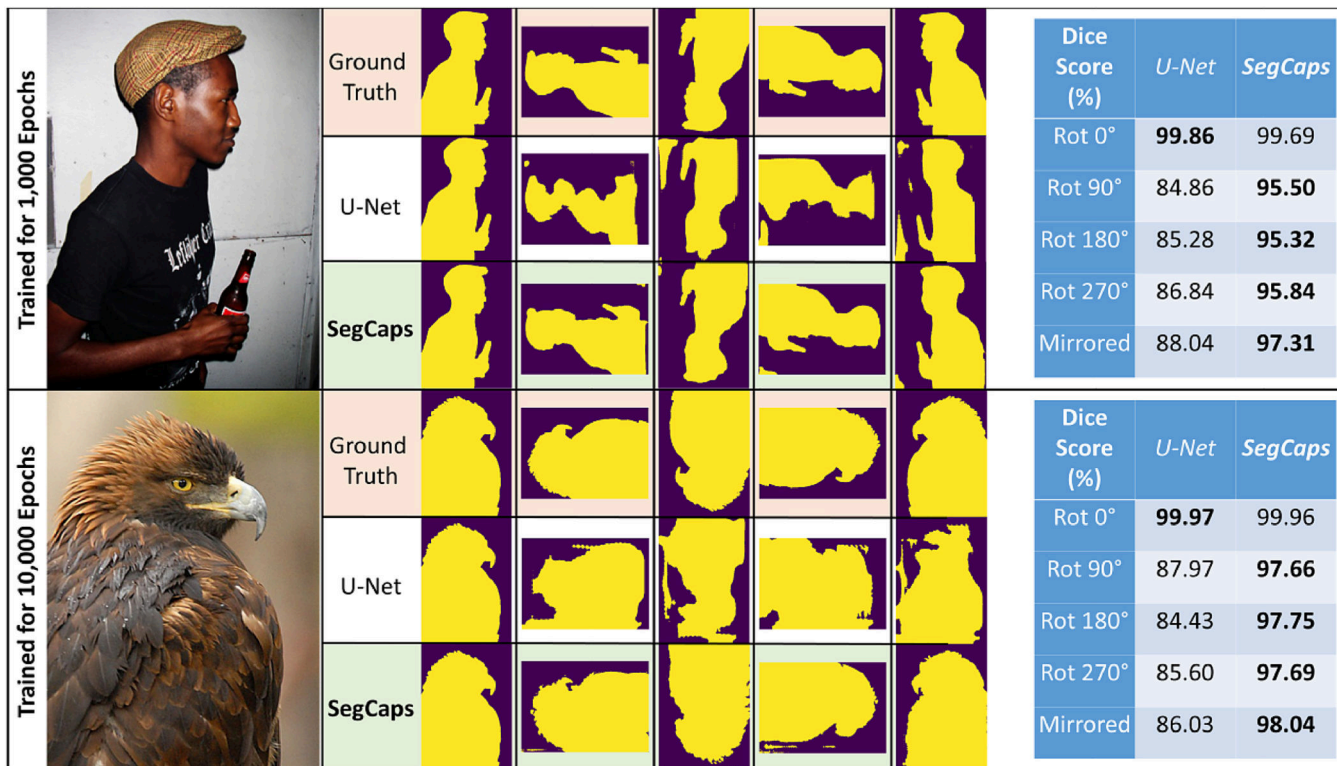


Figure 12: Testing the affine equivariant properties of capsule networks, specifically *SegCaps*, by overfitting on a single image, trained without augmentation, then predicting on transformations of that image.

Table 1:

Experimental results on 885 CT scans from the LIDC-IDRI database (Armato et al., 2011), measured by 3D Dice Similarity Coefficient and Hausdorff Distance (HD).

Method	Dice (% \pm std)	HD (mm \pm std)
<i>U-Net</i> (Ronneberger et al., 2015)	96.06 \pm 2.40	41.211 \pm 9.109
<i>Tiramisu</i> (Jégou et al., 2017)	94.40 \pm 3.66	42.205 \pm 15.210
<i>P-HNN</i> (Harrison et al., 2017)	95.64 \pm 2.92	41.775 \pm 13.866
<i>SegCaps</i>	96.98 \pm 0.36	30.764 \pm 2.793

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2:

Experimental results on 545 CT scans from the LTRC database (Karwoski et al., 2008), measured by 3D Dice Similarity Coefficient and Hausdorff Distance (HD).

Method	Dice (% \pm std)	HD (mm \pm std)
<i>U-Net</i> (Ronneberger et al., 2015)	95.52 \pm 2.80	37.625 \pm 6.831
<i>Tiramisu</i> (Jégou et al., 2017)	95.41 \pm 2.08	43.969 \pm 14.869
<i>P-HNN</i> (Harrison et al., 2017)	95.46 \pm 3.93	33.835 \pm 9.596
<i>SegCaps</i>	96.91 \pm 2.24	26.295 \pm 3.806

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 3:

Experimental results on 214 CT scans from the UHG database (Depeursinge et al., 2012), measured by 3D Dice Similarity Coefficient and Hausdorff Distance (HD).

Method	Dice (% \pm std)	HD (mm \pm std)
<i>U-Net</i> (Ronneberger et al., 2015)	88.10 \pm 1.84	44.303 \pm 34.148
<i>Tiramisu</i> (Jégou et al., 2017)	87.67 \pm 1.38	61.227 \pm 54.096
<i>P-HNN</i> (Harrison et al., 2017)	88.64 \pm 0.64	43.698 \pm 24.026
<i>SegCaps</i>	88.92 \pm 0.66	37.171 \pm 23.223

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 4:

Experimental results on 108 CT scans from the JHU-TBS database, measured by 3D Dice Similarity Coefficient and Hausdorff Distance (HD).

Method	Dice (% \pm std)	HD (mm \pm std)
<i>U-Net</i> (Ronneberger et al., 2015)	90.38 \pm 3.86	7.593 \pm 0.886
<i>Tiramisu</i> (Jégou et al., 2017)	86.45 \pm 5.76	7.428 \pm 1.337
<i>P-HNN</i> (Harrison et al., 2017)	88.81 \pm 6.81	7.517 \pm 1.896
<i>SegCaps</i>	93.35 \pm 0.95	4.367 \pm 1.367

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 5:

Experimental results on 208 CT scans from the JHU-TB database, measured by 3D Dice Similarity Coefficient and Hausdorff Distance (HD).

Method	Dice (% \pm std)	HD (mm \pm std)
<i>U-Net</i> (Ronneberger et al., 2015)	76.26 \pm 9.51	24.295 \pm 14.684
<i>Tiramisu</i> (Jégou et al., 2017)	79.99 \pm 6.24	24.647 \pm 11.629
<i>P-HNN</i> (Harrison et al., 2017)	80.11 \pm 7.46	26.597 \pm 16.168
<i>SegCaps</i>	80.91 \pm 5.27	26.021 \pm 10.260

Table 6:

Experimental results on water-only (fat-suppressed) MRI scans from the BLSA dataset, measured by 3D Dice Similarity Coefficient.

Method	Adipose (Fat) (% \pm std)	Muscle (% \pm std)
<i>FC</i> (Irmakci et al., 2018)	44.46 \pm 27.29	67.70 \pm 24.67
<i>U-Net</i> (Ronneberger et al., 2015)	84.48 \pm 8.33	90.00 \pm 1.85
<i>SegCaps</i>	84.45 \pm 6.60	90.74 \pm 1.49

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 7:

Experimental results on water-fat MRI scans from the BLSA dataset, measured by 3D Dice Similarity Coefficient.

Method	Adipose (Fat) (%± std)	Muscle (%± std)
<i>FC</i> (Irmakci et al., 2018)	78.52 ± 14.77	84.56 ± 14.66
<i>U-Net</i> (Ronneberger et al., 2015)	91.11 ± 3.10	90.34 ± 8.76
<i>SegCaps</i>	91.26 ± 2.77	92.59 ± 1.14

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 8:

Experimental results on fat-only (water suppressed) MRI scans from the BLSA dataset, measured by 3D Dice Similarity Coefficient.

Method	Adipose (Fat) (% \pm std)	Muscle (% \pm std)
<i>FC</i> (Irmakci et al., 2018)	77.52 \pm 16.38	73.00 \pm 17.78
<i>U-Net</i> (Ronneberger et al., 2015)	94.61 \pm 2.67	93.85 \pm 1.01
<i>SegCaps</i>	94.28 \pm 2.59	93.38 \pm 1.42

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 9:

Comparing the deeper encoder-decoder network structure *SegCaps* enabled by our proposed deconvolutional capsules, versus a network designed to be as similar as possible to CapsNet (Sabour et al., 2017) (*Baseline SegCaps*), abbreviated in table as Base-Caps.

Method	Dice (%± std)	HD (mm± std)
Base-Caps	75.97 ± 4.60 3	352.582 ± 133.451
<i>SegCaps</i>	88.92 ± 0.66	37.171 ± 23.223

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 10:

Number of parameters for each of the networks examined in this study. The percentage of less parameters (Percent Less) is measured relative to the number of parameters in *U-Net*.

Method	Parameters	Percent Less
<i>U-Net</i>	31.0 M	0.00 % Baseline (100 %)
<i>P-HNN</i>	14.7 M	52.58 47.42 %
<i>Tiramisu</i>	9.4 M	69.68 30.32 %
<i>Baseline SegCaps</i>	1.7 M	94.52 5.48 %
<i>SegCaps</i>	1.4 M	95.48 4.52 %

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 11:

Experimental results on the UHG dataset using downscaled version of *U-Net* and *Tiramisu* to roughly equal the same number of parameters (1.4 M) as *SegCaps*. The value of k (number of feature maps per layer reduction factor) for *U-Net* and *P-HNN* is included in parentheses.

Method	Dice (%± std)	HD (mm± std)
<i>U-Net</i> (orig.)	88.10 ± 1.84	44.303 ± 34.148
<i>U-Net</i> (4.68)	87.57 ± 2.80	62.006 ± 62.693
<i>Tiramisu-103</i>	87.67 ± 1.38	61.227 ± 54.096
<i>Tuamisu-56</i>	87.68 ± 0.96	67.913 ± 36.190
<i>P-HNN</i> (orig.)	88.64 ± 0.64	43.698 ± 24.026
<i>P-HNN</i> (3.2)	86.69 ± 1.39	82.223 ± 48.989
<i>SegCaps</i>	88.92 ± 0.66	37.171 ± 23.223

Table 12:

Examining the effect of the proposed extension of the reconstruction regularization to the task of segmentation.

Method	Dice (% \pm std)	hD (mm \pm std)
No Recon	88.58 \pm 1.03	42.345 \pm 21.180
With Recon	88.92 \pm 0.66	37.171 \pm 23.223

Table 13:

Examining the effect of different number of routing iterations (abbreviated as # ITERS) per forward pass of *SegCaps*. In 1,3, one routing iteration is performed when the spatial resolution remains the same and three iterations are performed when the resolution changes.

# ITERS	Dice (%± std)	HD (mm± std)
1	88.17 ± 1.23	67.668 ± 58.556
2	88.58 ± 1.03	42.345 ± 21.180
3	88.92 ± 0.66	37.171 ± 23.223
4	87.72 ± 1.36	110.901 ± 71.701
1,3	88.11 ± 1.13	72.877 ± 54.649

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript